Reliable Querying of Very Large, Fast Moving and Noisy Predicted Interaction Data using Hierarchical Crowd Curation

Hasan Jamil[♮] and **Fereidoon Sadri**[♭]

[♮]Department of Computer Science, University of Idaho, USA

jamil@uidaho.edu

[♭]Department of Computer Science, University of North Carolina at Greensboro, USA

f_sadri@uncg.edu

arXiv:1606.01957v1 [cs.DB] 6 Jun 2016

Abstract

The abundance of predicted and mined but uncertain biological data show huge needs for massive, efficient and scalable curation efforts. The human expertise warranted by any successful curation enterprize is often economically prohibitive especially for speculative end user queries that may not ultimately bear fruit. So the challenge remains in devising a low cost engine capable of delivering fast but tentative annotation and curation of a set of data items that can be authoritatively validated by experts later demanding significantly small investment. The aim thus is to make a large volume of predicted data available for use as early as possible with an acceptable degree of confidence in their accuracy while the curation continues. In this paper, we present a novel approach to annotation and curation of biological database contents using crowd computing. The technical contribution is in the identification and management of trust of mechanical turks, and support for ad hoc declarative queries, both of which are leveraged to support reliable analytics using noisy predicted interactions.

1 Introduction

Protein-protein interactions (PPI) are one of the central biological activities in almost all cellular processes. Thus, accurate structural and functional annotations of PPI are vital to understand their biological and pathological implications for new therapeutics design [19]. PPI network analyses are also often used to computationally predict protein functions to decipher molecular mechanisms and biological processes as an inexpensive alternative to more expensive laboratory methods [25]. PPI networks essentially provide a systems-level understanding of the underlying disease mechanisms, metabolic functions, help identify novel bio-markers, etc. [28]. Collection, curation and maintenance of such networks thus play a critical role in biological data analysis in today's post-genomic era.

To this end, researchers have been collecting interaction data to store in databases and annotate them for various applications for sometime [31, 22]. Integrating predicted and less accurate inter-

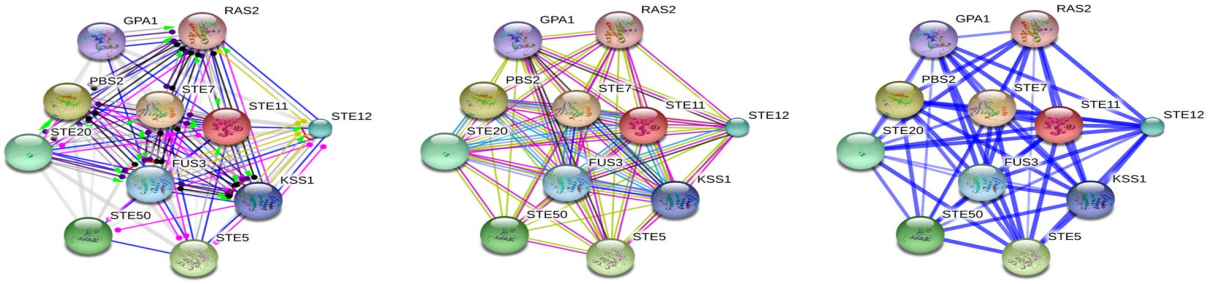

(a) Actions. (b) Evidence. (c) Confidence.

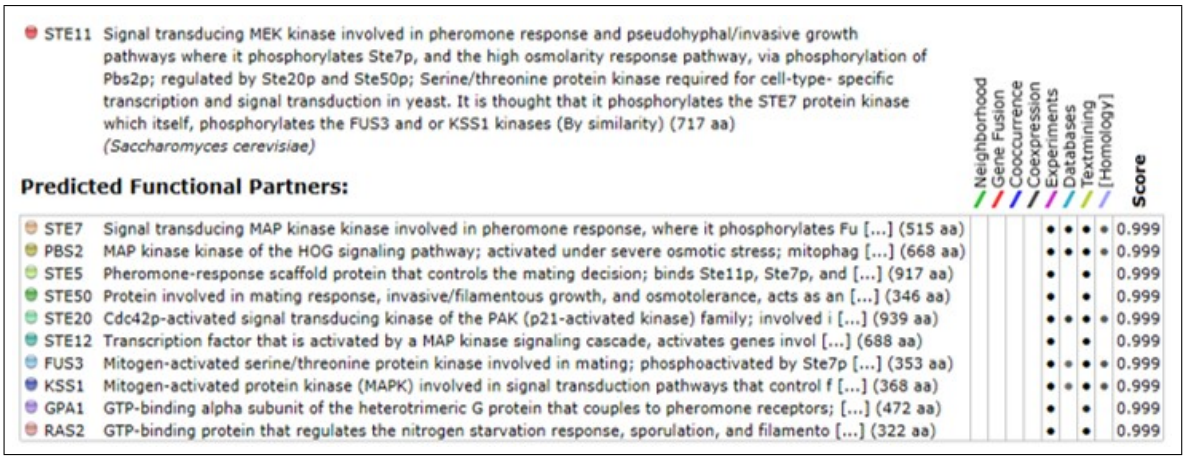
STE11 Signal transducing MEK kinase involved in pheromone response and pseudohyphal/invasive growth pathways where it phosphorylates Ste7p, and the high osmolarity response pathway, via phosphorylation of Pbs2p; regulated by Ste20p and Ste50p; Serine/threonine protein kinase required for cell-type- specific transcription and signal transduction in yeast. It is thought that it phosphorylates the STE7 protein kinase which itself, phosphorylates the FUS3 and or KSS1 kinases (By similarity) (717 aa)
(Saccharomyces cerevisiae)
Predicted Functional Partners:
Neighborhood
Gene Fusion
Cooccurrence
Coexpression
Experiments
Databases
Textmining
[Homology]
Score
STE7 Signal transducing MAP kinase kinase involved in pheromone response, where it phosphorylates Fu [...] (515 aa) 0.999
PBS2 MAP kinase kinase of the HOG signaling pathway; activated under severe osmotic stress; mitophag [...] (668 aa) 0.999
STE5 Pheromone-response scaffold protein that controls the mating decision; binds Ste11p, Ste7p, and [...] (917 aa) 0.999
STE50 Protein involved in mating response, invasive/filamentous growth, and osmotolerance, acts as an [...] (346 aa) 0.999
STE20 Cdc42p-activated signal transducing kinase of the PAK (p21-activated kinase) family; involved i [...] (939 aa) 0.999
STE12 Transcription factor that is activated by a MAP kinase signaling cascade, activates genes invol [...] (688 aa) 0.999
FUS3 Mitogen-activated serine/threonine protein kinase involved in mating; phosphoactivated by Ste7p [...] (353 aa) 0.999
KSS1 Mitogen-activated protein kinase (MAPK) involved in signal transduction pathways that control f [...] (368 aa) 0.999
GPA1 GTP-binding alpha subunit of the heterotrimeric G protein that couples to pheromone receptors; [...] (472 aa) 0.999
RAS2 GTP-binding protein that regulates the nitrogen starvation response, sporulation, and filamento [...] (322 aa) 0.999

(d) Overall STE11 prediction information for top ten interactors.

Figure 1: Various interaction graph views of protein STE11.

action data with more reliable laboratory tested data has been an approach essentially to increase the coverage of the network analyses for novel discoveries [34]. To expedite discovery, researchers have been increasingly adopting more predicted interactions [29] into their analyses obtained from high throughput data generation systems such as next generation sequencing [30], yeast two-hybrid screens [24, 18] mass spectrometry [32] and literature mining [9]. Unfortunately the high error rates in these predicted data sets invite two major challenges. First, use of these datasets render the analysis highly speculative and often wrong. Secondly, until these datasets are curated [5] beyond an acceptable level of certainty, they essentially remain less than useful.

Manual or machine curation of these predicted interactions is often a monumental task. Additionally, autonomous [2], semi-autonomous [10] and manual curation [14] have varying degrees of accuracy rates along with their respective scale factors. The BioCreative II effort [13] recently studied the difficulties associated with the curation and annotation process of PPI data. They found that only about 35% precision and recall could be achieved for mapping interactions obtained from text mining to SwissProt, and only about 20% of predicted interactions from full-text article mining could be matched with human curator identified sentences. Such low success rates reflect the

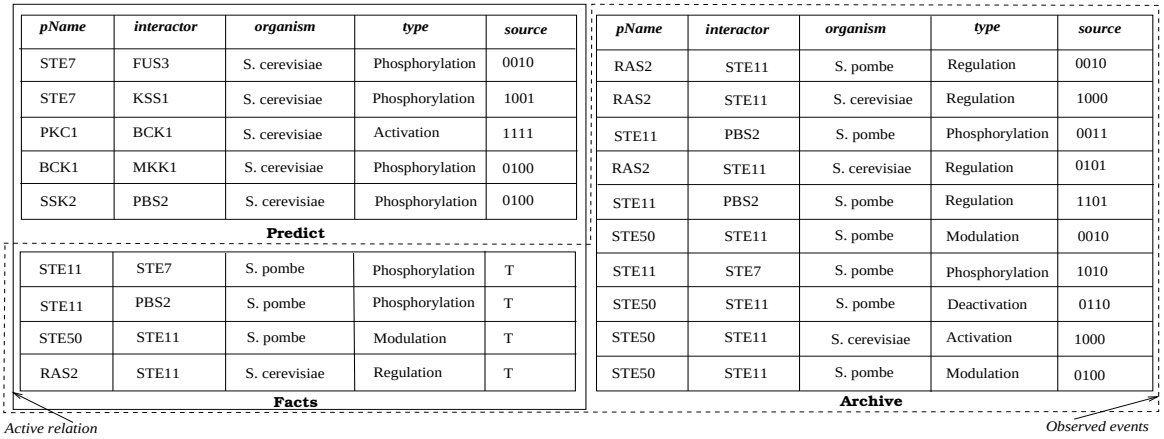
pName interactor organism type source
STE7 FUS3 S. cerevisiae Phosphorylation 0010
STE7 KSS1 S. cerevisiae Phosphorylation 1001
PKC1 BCK1 S. cerevisiae Activation 1111
BCK1 MKK1 S. cerevisiae Phosphorylation 0100
SSK2 PBS2 S. cerevisiae Phosphorylation 0100
Predict
STE11 STE7 S. pombe Phosphorylation T
STE11 PBS2 S. pombe Phosphorylation T
STE50 STE11 S. pombe Modulation T
RAS2 STE11 S. cerevisiae Regulation T
Facts
pName interactor organism type source
RAS2 STE11 S. pombe Regulation 0010
RAS2 STE11 S. cerevisiae Regulation 1000
STE11 PBS2 S. pombe Phosphorylation 0011
RAS2 STE11 S. cerevisiae Regulation 0101
STE11 PBS2 S. pombe Regulation 1101
STE50 STE11 S. pombe Modulation 0010
STE11 STE7 S. pombe Phosphorylation 1010
STE50 STE11 S. pombe Deactivation 0110
STE50 STE11 S. cerevisiae Activation 1000
STE50 STE11 S. pombe Modulation 0100
Archive
Active relation
Observed events

Figure 2: An eIST interaction database **D**.

difficulty of obtaining meaningful and accurate curation, and the scopes that exist to design novel ways to curate or use the uncurated interactions. The need is so pressing that it motivated the desperate attempt to even manually curate about 88,000 documents [7]. In this paper, our goal is addressing this niche and propose a method to make uncurated interaction data meaningfully usable in an evolving database.

1.1 An Illustrative Example using STRING and eIST

Consider extracting all interactors of the Serine/threonine-protein kinase STE11 in Saccharomyces cerevisiae from the interaction database STRING [31]. STRING supports multiple types of queries to search all known and predicted interactors of a given set of proteins or gene products. Figures 1(a) through 1(c) show three views of ten most prominent interactors of the STE11 in a single protein query mode and their summary characterizations are shown in figure 1(d).

We can imagine that the three views of the interacting partners of STE11 in figures 1(a) through 1(c) are likely generated from a traditional relational tabular view of the interactions functionally similar to the extended information source tracking (eIST) [11] tables as shown in figure 2. In this earlier research on eIST, all relations are composed of three complementary tables having identical schema – *Predict, Facts*, and *Archive*. The tuples in *Facts* are regular relational tuples, except that all eIST tables have an additional system recognized attribute called *source* which contains a vector of sources contributing the tuple.

Assuming that we have four contributing curators, or sources, in this example database – *Russ, Fred, Karen* and *Maria* in this order, and of course the special system source `T`, the tuple `<STE11,STE7,S. pombe,Phosphorylation@1010>` in *Archive* means that *Russ* and *Karen* con-

tributed this tuple, and the validity and reliability of it now depends on the combined reliability of both contributors. However, its presence in *Archive* signifies that it is an *image* of a `true` tuple in *Facts* and thus does not take part in any query. Tuples in *Archive* migrate from *Predict* every time a `true` tuple corresponding to it is added in the *Facts* table, and are only used to determine the credibility of the sources based on their prediction or curation performance. For example, if we add `<BCK1,MKK1, S. cerevisiae,Activation@T>` in *Facts*, the tuple `<BCK1,MKK1,S. cerevisiae, Phosphorylation@0100>` in *Predict* will migrate to *Archive* since the primary key components of both tuples (i.e., `<BCK1,MKK1, S. cerevisiae>` agree, and once we know that `<BCK1,MKK1, S. cerevisiae,Activation>` is a fact (i.e., 100% reliable), there is no reason for `<BCK1,MKK1,S. cerevisiae, Phosphorylation @0100>` to continue to exist in *Predict*. The source vector `T` is a special system source which is always correct, and thus is the $n+1$th source. Therefore, for the tuple `<BCK1,MKK1,S. cerevisiae,Phosphorylation @0100>`, vector `<0100>` in reality is `<01000>` and for tuple `<BCK1,MKK1, S. cerevisiae, Activation@T>`, it is `<00001>`, with the proviso that whenever $n+1$ bit of a vector is 1, the rest must be zeros, it is 0 otherwise.

The basic IST model was introduced earlier in [27], and the algebraic model and reliability calculation method of eIST has been later devised by Jamil and Sadri in [11] to accommodate the time varying nature of source reliability based on past performance. In these two works, it had been shown that these models are capable of associating a confidence score to any query tuple computed using arbitrary `Select Project Join` (SPJ) query based on eIST relations. The eIST model includes source vector manipulations for all basic relational algebra operations, and a reliability calculation algorithm based on source credibility. These models provide the ability to determine the *provenance* of each query tuple. That is, how the tuple was obtained and which sources contributed. In itself, it is a valuable tool for explaining how the answer was derived. In addition, a probability or certainty factor can be computed for each query tuple using the provenance. IST and eIST use solid mathematical probability theory for the probability calculation. But it is easy to accommodate other approaches [15], such as the fuzzy logic technique, for the computation of certainty factors. We omit a detailed discussion of these models here for brevity and refer the readers to [11, 27] for an in-depth discussion.

Example 1.1 (Calculating Tuple Reliability) *Consider the tuple* `<BCK1, MKK1, S. cerevisiae, Phosphorylation@0100>` *in* Predict*, which shows that source 2,* Fred*, contributed to this information. The reliability of sources are maintained and updated periodically in eIST. Assume the reliability of* Fred *has been determined to be 0.8. Then the certainty factor of the tuple* `<BCK1,`

MKK1, S. cerevisiae, Phosphorylation@0100> *is 0.8, or 80%. Now assume that this information was contributed by* Fred *and by* Karen*, which would be represented by tuple* <BCK1,MKK1, S. cerevisiae, Phosphorylation@0110> *in* Predict*. If reliability of* Karen *has been determined to be 0.85, we expect the certainty of this tuple to be very high, since these sources are* independently *confirming this information. The certainty factor of* <BCK1,MKK1, S. cerevisiae, Phosphorylation@0110> *is computed as* $1 - (1 - 0.85) \times (1 - 0.8) = 0.97$ *in the probabilistic approach. Note that in the fuzzy logic approach, the certainty factor of this tuple is computed as the higher reliability of contributing sources, namely, 0.85. Fuzzy logic formulation may be relevant in some applications, and eIST can accommodate either approach depending on users preference.*

However, the observation we would like to make here is that the three views, i.e., actions, evidence and confidence, that STRING generates can be easily constructed from an eIST database with the added benefit of a time varying confidence score in a dynamic fashion. The partitioning of the tuples in *Predict* and *Facts* also makes it possible to start using raw data without serious expert curation as soon as they become available in *Predict* and need not be in *Facts*. The eight evidential attributes in figure 1(b) STRING uses to assign a confidence score in figure 1(c) can each be modeled as a source and their combined reliability can be used to generate a network similar to figure 1(c) in eIST. We leverage this observation in this paper to propose a new crowdsourced curation model for PPI networks using a declarative query language, called *CureQL*.

1.2 Using CrowdDB for Content Gathering from People

Leveraging crowd participation for content gathering is the principal focus of databases such as CrowdDB [8] and Qurk [16]. They support a declarative language to initiate crowd participation for specific data items or human observations such as subjective comparison of images. They also support several ancillary functions to manage the collection process at the systems level including language support for managing Human Intelligent Task (HIT) and HIT assignments, active time of the crowd queries, and query specific fabrication of the query interfaces. We adapt CrowdSQL (see [8] for more details) supported in CrowdDB as a substrate for CureQL to inherit its inherent simplicity and leverage its powerful features.

In the examples in figure 3, the table in c_1 is called a *crowd table*, and the table in c_2 is a table with *crowd columns*, but both are called *crowd enabled tables*. The semantics of the CrowdSQL's [8] crowd enabled tables are straightforward – when a crowd enabled relation is referenced in an query, entire rows are extracted with crowd participation corresponding to a crowd table as opposed

c_1: CREATE CROWD TABLE *Interaction* (
pName STRING,
interactor STRING,
organism STRING,
type STRING,
PRIMARY KEY (*pName, interactor, organism*),
FOREIGN KEY (*pName, interactor, organism*),
REF *Names* (*name, name, organism*));

c_2: CREATE TABLE *Names* (
name CROWD STRING,
organism STRING,
date CROWD DATE,
PRIMARY KEY (*name,*
organism));

q_1: SELECT *pName, organism*
FROM *Interaction*
WHERE *type = "activation"*;

q_2: SELECT *name, organism*
FROM *Names*
WHERE *date* $\geq$ *"1/31/2016"*;

Figure 3: Crowd specification of database **D** in figure 2 in CrowdSQL/CureQL.

to columns extracted corresponding to a table with crowd columns. For example, the query q_1 in figure 3 asks the crowd to supply entire rows and add to the existing *Interaction* table since the entire table is crowd enabled. However, the query q_2 asks for only the the missing values (identified with the presence of a CNULL, or crowd NULL value) for the columns *name* and *date*. Both queries, however, will finally select only the tuples that satisfy the corresponding selection conditions.

1.3 CrowdCure Approach and Contributions

While the CrowdDB and eIST frameworks individually support some of the complementary features we envision in a community curation system, several additional features are still necessary to develop an end to end working system, we call *CrowdCure*. For example, support for the use of specific extraction tools on target data sources, assignment of desired and relevant curators to sets of predicted data, organizing the curators in an expertise hierarchy, etc. within a declarative setting to aid ad hoc queries.

From an end user standpoint, we expect a user to be able to pose both ad hoc traditional eIST, and crowd curation queries in CrowdCure. By that we mean, when new contributions from crowd aren't necessary, we should be able to pose simple queries and get traditional responses with tuple reliability. Recall that all tables in eIST are potentially uncertain (their certainty ranges between the interval (0,1]) and their certainty is determined by the reliability of the contributing sources, and that all tables have a special attribute called *source* and all tables have three possibly empty

components *Facts, Predict* and *Archive*.

We, therefore, adopt the provision of explicit invocation of crowd curation features even on crowd enabled tables. We believe this approach is pragmatic and ensures flexible use of the underlying database to suit the application semantics as needed. Thus from a systems standpoint, we support an interface in which users can issue any mix of extraction, curation and analysis queries directly using a declarative language. In extraction queries, linguistic features are supported to specify mining tools, set of documents, and to partition extraction tasks for groups of curators organized in a hierarchy. Visual interfaces are supported for each of these tasks. For analysis queries, the traditional features of eIST becomes handy. The approach also supports separation of extraction from a hierarchical curation, and as soon as the data are available for curation, they can be used in queries. As the curation progresses, their reliability improves (positively or negatively) and are reflected in the composite SPJ query response in a dynamic temporal fashion. Finally, another desirable feature of CrowdCure is the ability of eIST to display the current reliabilities it has determined for all sources. We allow users to override eIST reliabilities and supply their own reliabilities for one or more sources if they wish. We also support, among other features, "what-if" queries allowing users to experiment with the effects of changing source reliabilities.

2 Related Research

In CrowdCure, we are dealing with research challenges in two orthogonal axes – crowd curation, and assignment of reliability to curated interaction data. A related issue is how to quantify trust of computed results of arbitrary queries using the base inaccurate data curated by numerous potentially unreliable experts. The challenge is to marry these two orthogonal advances into one single coherent platform for reliable crowd curation and supporting queries with quantifiable reliability in a user transparent manner.

Collaborative or crowd computing rely on social participation of masses toward a common goal such as content curation. Two most widely employed curation types are expert-based annotation and machine generated classification. Traditionally, however, the curation was solely an expert's job. Algorithmic curation, on the other hand, ranks curation items such as relevant web pages for search queries or produce recommendations on e-commerce databases for user consumptions. In contrast the voting based annotation commonly employed in e-commerce applications is high in popularity due to its simplicity and amenability to natural ranking [3] because the upvote and downvote based ranking of curation overcomes a few drawbacks of expert-based curation (cost

of experts and scale) and algorithm based curation (inflexibility of algorithms of content types). However, the voting based curation has not been shown to be effective in technical content ranking which demand their understanding scientifically. Technically, interaction annotation require human understanding of an interaction pair in a document that is potentially confusing and fundamentally distinct from an overall likability of an entire document or a piece of item. In other words, technical curation usually involves extraction of structured data with various attributes from texts to be stored in a structured database such as relational database.

q_3: SELECT *pName, interactor, type*
USING *UniHi* ON *I_papers*
SOURCE r_1 BEFORE SELECT *eName*
FROM *Experts*
WHERE *laboratory = "MIT"*;
FROM *Interaction*
WHERE *type = "activation"*;

q_4: SELECT *t.pName, t.interactor, t.type*
USING *UniHi* ON *I_papers*
LIMIT TIME t UNITS, DATA k ROWS
FROM *Interaction* AS y
GROUP BY *type*
CLUSTER SOURCE (SELECT *eName, laboratory*
FROM *Experts*
WHERE *y.organism = species*) BEFORE
(SELECT *eName, species*
FROM *Experts*);

c_3: CREATE TABLE *Experts* (
eName STRING,
species STRING,
laboratory STRING,
eMail STRING UNIQUE,
SOURCE KEY (*eName*),
PRIMARY KEY (*eName, species*));

c_4: CREATE VIEW r_1 AS
SELECT *eName*
FROM *Experts*
WHERE *laboratory = "Yale"*;

q_5: SELECT *name, organism*
USING *PIPs* ON *P_papers*
SOURCE r_2 BEFORE r_3
FROM *Names*;

Figure 4: CureQL queries over the database **D** (continued from figure 3).

Research in collecting structured data has gained its due attention recently [23]. In these research, turks are asked to identify structured data in texts in order to populate structured tables in ways similar to extracting interaction data from free texts [29]. The process of extracting interaction data by turks is facilitated by suitable mining and text highlighting interfaces in some research to help extract quality data [26, 21]. Several research also tried to quantify the quality or reliability of extracted data in some fashion [1], and several adopted an alternative method of algorithmic validation of extracted data [10, 2]. An entirely different approach avoids assignment

of reliability to or validation of extracted data by computationally choosing reliable turks in the first place [4]. Our understanding is that none of these approaches scale well for large data sets, exhibit unreliable certainty assignment and skewed quantification [33].

3 CureQL Language and its Semantics

Technically, CureQL extends CrowdSQL in two ways: as shown in figure 4, (i) it introduces the notion of `SOURCE KEY` in the `CREATE TABLE` and `CREATE CROWD TABLE` statements to declare a set of attributes to be identifiers a la object-oriented systems to be used as curator identities, and (ii) it introduces an optional `USING ON` clause with an optional `LIMIT` and a mandatory `SOURCE` subclause, with the proviso that the `SOURCE` subclause may be used with the `USING` clause globally, or with the `GROUP BY` clause for every partition locally.

Consider the use of `USING ON` clause and the `SOURCE` subclause in query q_3 in figure 4. `USING` subclause has no effect on a query if it does not include a crowd column or the table is not a crowd table. When used, it expects an analytics, such as *UniHi* [12] (query q_5 uses *PIPs* [17], *ChiliBot* [6] is another example), which requires an input *I_papers*. The `ON` modifier is optional to allow analytics without input needs. However, the analytics must return all the crowd columns for the `SELECT` query to be successful. Since `USING` is an extraction operation, and all extracted data possibly are predictions (all crowd contributions are treated as uncertain), all tuples now must be tagged with a set of contributing sources or curators. In CrowdQL, the first set of contributors are also treated as curators. This tagging is accomplished using the mandatory `SOURCE` subclause.

The `SOURCE BEFORE` subclause expects an arbitrary but totally ordered sets of source keys to serve as the contributing or curating experts of the predicted data. Each scheme R_i of relations r_i in the subclause of the form `SOURCE` r_1 `BEFORE ... BEFORE` r_k must contain a unique `SOURCE KEY`. The projection of the `SOURCE KEY`s are the assigned curators for each extracted tuple in the *Predict* set. However, the set of curators in r_j are strictly at a higher level than the curators in set r_i for any $i < j$ and are more credible. As query q_3 shows, CureQL supports supplying these sets as either a ground relation, a view (as in the statement c_4) or as a CureQL query (in queries q_3 and q_4) satisfying the `SOURCE KEY` requirement. In query q_5, if none of the attributes of r_2 (or r_3) is part of a `SOURCE KEY`, the entire scheme is treated as one.

The scope rule of `SOURCE BEFORE` subclause can be used in two interesting ways. First, the relations in the `FROM` clause in query q_3 do not range over the `SOURCE BEFORE` subclause as it appeared before the `FROM` clause, and thus the use of tuple variables in the CureQL queries in

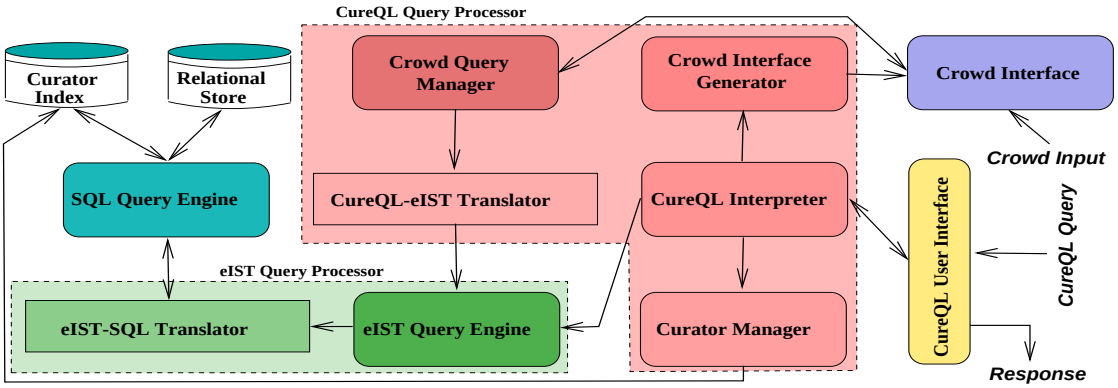

Figure 5: CrowdCure database system architecture.

`SOURCE BEFORE` subclause has no effect. Second, in the query q_4, the relation *Interaction* has scope over the `SOURCE BEFORE` subclause that follows the `FROM` clause and thus influences the selection of the experts in `SOURCE` subclause based on the species value for the first curator set in the order. Thus, the the use of `CLUSTER` modifier before the `SOURCE BEFORE` subclause in `GROUP BY` clause forces assigning potentially different sets of curator hierarchies for every group partitioned by the `GROUP BY` statement. However, note that the `SOURCE` clause cannot be used in both `USING` and `GROUP BY` clauses simultaneously. Finally, in order to deal with uncertainties of long crowd response delay and potentially large volumes, CureQL supports the `LIMIT` option slightly differently than Qurk [16]. CureQL query q_4 specifies a time limit of t units, and a maximum number of k rows although, in principle, a crowd query posed can be indefinite aiming to collect rows perpetually.

4 CrowdCure System Architecture

The overall architecture of CrowdCure system is shown in figure 5. In this architecture, we leverage an eIST query processor, which in turn is built as a front-end to a relational query engine, and a relational storage system. Essentially this means that all eIST tables and queries are mapped onto a relational schema for the SQL engine to manage and process. The source vector manipulation and reliability assignment component of a query are managed by the the eIST engine orthogonally by using a separate system index for source information called the *Curator Index*. User queries are submitted via the *CureQL User Interface* to the *eIST Query Engine* transparently of the back-end eIST system after they are preprocessed using the *CureQL Interpreter* and translated using the *CureQL-eIST Translator*. We briefly describe the *CureQL Query Processor* implementation in the next few sections that are unique to CrowdCure system.

4.1 CureQL Query Processor Implementation

To process CureQL queries, we implement a minimally invasive stack of two native query processors – the *eIST Query Processor* and the *SQL Engine*. To avoid the need for a complete overhaul of eIST, we design an interface between eIST and CureQL, called the *CureQL-eIST Translator*, to map all CureQL queries into an equivalent eIST query involving crowd enabled relations. To achieve the effect of the CrowdSQL and Qurk fusion, we have added three subsystems called the *Crowd Interface Generator*, *Crowd Query Manager* and *Curator Manager* to capture the unique ways the `USING ON` and `SOURCE BEFORE` subclauses work. Once a query is received by the *CureQL Interpreter*, it is transmitted directly to the eIST engine if no crowd columns are involved in the query scheme. Otherwise, for each table in the `FROM` clause that involves a crowd column, appropriate crowd input forms are created and floated on the web to facilitate curator inputs by the *Crowd Interface Generator* unit.

The *Crowd Interface Generator* translates `SOURCE` tables into authorization rules and access rights for the curators. For example, if *Fred* and *Maria* are the curators for a crowd tuple t, they will see a request for input in their system while *Karen* and *Russ* will not. `GROUP BY CLUSTER SOURCE` also works in an analogous way. If *Karen* and *Russ* are at a higher level in the curator hierarchy, we technically have two choices. In the first case, they will see tuple t once *Fred* and *Maria* have submitted their input, or their time has expired, and thus allowing *Russ* and *Karen* to essentially cure the first curation. The other option is to let *Karen* and *Russ* input their values at anytime and calculate a combined credibility giving higher weight to *Russ* and *Karen*'s inputs as they are at a strictly higher curation stratum. We adopted the first choice in the current version of CrowdCure system. Thus, depending on the use of the `SOURCE BEFORE` placement in the query, partition information and curator hierarchies are maintained for the entire duration of the query in coordination with the *Curator Manager* and *Curator Index*. Finally, the *Crowd Query Manager* handles the `LIMIT` options along with multi-table queries, `CLUSTER SOURCE`, curator hierarchy, and `SOURCE KEY` consistency enforcement in `SOURCE BEFORE` clause.

4.2 Query Processing

Since all tables in CrowdCure are essentially eIST tables (potentially with no *Predict* or *Archive* components), no queries are processed directly by the *SQL Engine*, and so we are able to classify a query into a purely eIST query (with no crowd columns in any tables) or a CureQL query (with having at least one crowd column). Since eventually all queries are translated into eIST queries by

the *CureQL-eIST Translator* or are directly processed as eIST queries, there is no technical hurdle in preserving the intended semantics of CureQL through staged translation approach adopted. The reliability calculation of the tuples in the query view utilizes the reliability mapping algorithm based on the source credibility, which in turn depends on the past history of curation of the sources as outlined for the basic IST [27] and eIST [11] frameworks, as a final step. For the current version, we have adopted a probabilistic model for reliability calculation, and different or even a parametric model [15] is equally possible.

5 Summary

Our goal in this paper was to introduce a novel hierarchical crowd curation system for fast moving, unreliable and very large PPI data sets. The motivation was to make the predicted data immediately available for use by the end users even before they are decisively cured. Based on multiple prototype implementations and trial runs, we have reasons to believe that this approach delivers more accurate curation with less effort. While limited types of queries are supported in the current prototypes, they supported ad hoc querying in a flexible way through a QBE type graphical interface. A full implementation of CrowdCure system is currently underway. Our ultimate goal is to use STRING [31] and DroID [20] as test cases and compare them with CrowdCure in terms of usability, accuracy, scale, curation delay or response time, and expressibility as a future research. This research also opens up a new dimension in crowd computing and novel HIT management interface design that we also plan to pursue.

References

[1] O. Alonso, C. C. Marshall, and M. A. Najork. A human-centered framework for ensuring reliability on crowdsourced labeling tasks. In *Human Computation and Crowdsourcing: Works in Progress and Demonstration Abstracts, An Adjunct to the Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, November 7-9, Palm Springs, CA, USA*, 2013.

[2] A. Antony, S. Basetty, S. Hartanto, and M. J. Palakal. Computational approach to biological validation of protein-protein interactions discovered using literature mining. In *ACM Symposium on Applied Computing (SAC), Fortaleza, Ceara, Brazil, March 16-20*, pages 1302–1306, 2008.

[3] G. Askalidis and G. Stoddard. A theoretical analysis of crowdsourced content curation. In *Workshop on Social Computing and User Generated Content*, 2013.

[4] A. Bozzon, M. Brambilla, S. Ceri, M. Silvestri, and G. Vesci. Choosing the right crowd: expert finding in social networks. In *Joint EDBT/ICDT Conferences, EDBT Proceedings, Genoa, Italy, March 18-22*, pages 637–648, 2013.

[5] D. Cao, N. Xiao, Q. Xu, and A. F. Chen. Rcpi: R/bioconductor package to generate various descriptors of proteins, compounds and their interactions. *Bioinformatics*, 31(2):279–281, 2015.

[6] H. Chen and B. M. Sharp. Content-rich biological network constructed by mining pubmed abstracts. *BMC Bioinformatics*, 5(1):1–13, 2004.

[7] A. P. Davis, T. C. Wiegers, P. M. Roberts, B. L. King, J. M. Lay, K. Lennon-Hopkins, D. Sciaky, R. J. Johnson, H. Keating, N. Greene, R. Hernandez, K. J. McConnell, A. Enayetallah, and C. J. Mattingly. A ctd-pfizer collaboration: manual curation of 88,000 scientific articles text mined for drug-disease and drug-phenotype interactions. *Database*, 2013, 2013.

[8] M. J. Franklin, D. Kossmann, T. Kraska, S. Ramesh, and R. Xin. Crowddb: answering queries with crowdsourcing. In *ACM SIGMOD International Conference on Management of Data, Athens, Greece, June 12-16*, pages 61–72, 2011.

[9] S. Haider, Z. Lipinszki, M. Przewloka, Y. Ladak, P. D'Avino, Y. Kimata, P. Liò, and D. Glover. DAPPER: a data-mining resource for protein-protein interactions. *BioData Mining*, 8:30, 2015.

[10] D. G. Jamieson, M. Gerner, F. Sarafraz, G. Nenadic, and D. L. Robertson. Towards semi-automated curation: using text mining to recreate the hiv-1, human protein interaction database. *Database*, 2012, 2012.

[11] H. M. Jamil and F. Sadri. Recognizing credible experts in inaccurate databases. In *ISMIS*, pages 46–55, 1994.

[12] R. K. R. Kalathur, J. P. Pinto, M. A. Hernndez-Prieto, R. S. Machado, D. Almeida, G. Chaurasia, and M. E. Futschik. Unihi 7: an enhanced database for retrieval and interactive analysis of human molecular interaction networks. *NAR*, 2013.

[13] M. Krallinger, F. Leitner, C. Rodriguez-Penagos, and A. Valencia. Overview of the protein-protein interaction annotation extraction task of BioCreative II. *Genome biology*, 9 Suppl 2(Suppl 2):S4+, 2008.

[14] D. Kwon, S. Kim, S. Shin, A. Chatr-aryamontri, and W. J. Wilbur. Assisting manual literature curation for protein-protein interactions using bioqrator. *Database*, 2014, 2014.

[15] L. V. S. Lakshmanan and N. Shiri. A parametric approach to deductive databases with uncertainty. *IEEE Trans. Knowl. Data Eng.*, 13(4):554–570, 2001.

[16] A. Marcus, E. Wu, S. Madden, and R. C. Miller. Crowdsourced databases: Query processing with people. In *Biennial Innovative Data Systems Research Conference, Asilomar, CA, USA, January 9-12, Online Proceedings*, pages 211–214, 2011.

[17] M. D. McDowall, M. S. Scott, and G. J. Barton. PIPs: human protein-protein interaction prediction database. *NAR*, 37(suppl 1):D651–D656, 2009.

[18] J. Mehla, J. H. Caufield, and P. Uetz. Mapping protein-protein interactions using yeast two-hybrid assays. *Cold Spring Harbor Protocols*, 2015(5), 2015.

[19] I. H. Moal, B. Jiménez-García, and J. Fernández-Recio. Ccharppi web server: computational characterization of protein-protein interactions from structure. *Bioinformatics*, 31(1):123–125, 2015.

[20] T. Murali, S. Pacifico, J. Yu, S. Guest, G. G. Roberts, and R. L. Finley. DroID 2011: a comprehensive, integrated resource for protein, transcription factor, RNA and gene interactions for drosophila. *Nucleic Acids Research*, 39(suppl 1):D736–D743, Jan. 2011.

[21] R. T. Nakatsu and C. L. Iacovou. An investigation of user interface features of crowdsourcing applications. In *1st International Conference on HCI in Business, Held as Part of HCI International, Crete, Greece, June 22-27*, pages 410–418, 2014.

[22] S. E. Orchard, M. Ammari, B. Aranda, L. Breuza, and et al. The MIntAct project - IntAct as a common curation platform for 11 molecular interaction databases. *Nucleic Acids Research*, 42(Database-Issue):358–363, 2014.

[23] H. Park and J. Widom. Crowdfill: collecting structured data from the crowd. In *ACM SIGMOD, Snowbird, UT, USA, June 22-27*, pages 577–588, 2014.

[24] J. R. Parrish, K. D. Gulyas, and R. L. Finley. Yeast two-hybrid contributions to interactome mapping. *Current opinion in biotechnology*, 17(4):387–393, Aug. 2006.

[25] W. Peng, J. Wang, J. Cai, L. Chen, M. Li, and F.-X. Wu. Improving protein function prediction using domain and protein complexes in ppi networks. *BMC Systems Biology*, 8(1):1–13, 2014.

[26] B. Rahmanian and J. G. Davis. User interface design for crowdsourcing systems. In *International Working Conference on Advanced Visual Interfaces, Como, Italy, May 27-29*, pages 405–408, 2014.

[27] F. Sadri. Reliability of answers to queries in relational databases. *IEEE Trans. Knowl. Data Eng.*, 3(2):245–251, 1991.

[28] T. Sevimoglu and K. Y. Arga. The role of protein interaction networks in systems biomedicine. *Computational and Structural Biotechnology J.*, 11(18):22 – 27, 2014.

[29] S. Subramani, R. Kalpana, P. M. Monickaraj, and J. Natarajan. Hpiminer: A text mining system for building and visualizing human protein interaction networks and pathways. *Journal of Biomedical Informatics*, 54:121–131, 2015.

[30] B. Suter, X. Zhang, C. G. Pesce, A. R. Mendelsohn, S. P. Dinesh-Kumar, and J.-H. H. Mao. Next-Generation sequencing for binary Protein-Protein interactions. *Frontiers in genetics*, 6, 2015.

[31] D. Szklarczyk, A. Franceschini, S. Wyder, K. Forslund, D. Heller, J. Huerta-Cepas, M. Simonovic, A. Roth, A. Santos, K. P. Tsafou, M. Kuhn, P. Bork, L. J. Jensen, and C. von Mering. String v10: Protein-protein interaction networks, integrated over the tree of life. *Nucleic Acids Research*, 43(D1):D447–D452, 2015.

[32] J. Vasilescu and D. Figeys. Mapping protein-protein interactions by mass spectrometry. *Current Opinion in Biotechnology*, 17(4):394–399, Aug. 2006.

[33] S. Xie, Q. Hu, J. Zhang, J. Gao, W. Fan, and P. S. Yu. Robust crowd bias correction via dual knowledge transfer from multiple overlapping sources. In *IEEE Intl. Conf. on Big Data, CA, USA, October 29 - November 1*, pages 815–820, 2015.

[34] Y. Zhang, H. Lin, Z. Yang, and J. Wang. Integrating experimental and literature protein-protein interaction data for protein complex prediction. *BMC Genomics*, 16(S-2):S4, 2015.